HP OpenView

Storage Mirroring application notes

Guidelines for using Pervasive SQL with Storage Mirroring

Part number: T2558–88058
Second edition: June 2006

Legal and notice information

Guidelines for using Pervasive SQL with Storage Mirroring application notes

IDocument Overview

This document is a *Storage Mirroring* application note. An application note provides guidelines on the use of Storage Mirroring in a specific environment.

This document contains:

- **Document Overview**—Explains what an application note contains, how it should be used, what you need to know before trying to use the application note, and where you can go for more information.
- **Solution Overview**—Explains how the application works with Storage Mirroring and describes the considerations that you must weigh when implementing your Storage Mirroring solution. Review this section to make sure that you understand the theory involved with using Storage Mirroring and your application. Includes both basics, such as system requirements, as well as configuration and environment-specific topics, such as interactions with specific clients or special considerations for WAN (Wide Area Network) environments. Pay special attention to those topics that are directly related to your environment.

Audience

This document is written for network and application administrators who have a working understanding of the applications and environments where the Storage Mirroring solution is to be deployed. You many need to expand on the documented information in order to customize the solution to fit your environment.

Before you use this application note, you should have an understanding of:

- Storage Mirroring
- Pervasive SQL

Expectations

Application notes are intended to provide a framework for configuring a Storage Mirroring solution in a specific environment and to draw attention to decisions you will need to make when configuring your solution.

Because there are an infinite number of possible configuration, network, and environment scenarios, application notes contain general configuration guidelines as well as an example configuration procedure that has been tested for a specific environment.

This document assumes that you are comfortable working with your operating system, Storage Mirroring, and Pervasive SQL.

Related documentation

Before you begin to configure your solution, make sure that you have complete documentation for your operating system, application, and Storage Mirroring. This application note does not provide step-by-step instructions for using standard operating system, application, and Storage Mirroring functionality.

The following documents contain additional information that you may need while setting up this solution:

- *HP OpenView Storage Mirroring user's guide* or online documentation
- Reference guides or documentation for Pervasive SQL

Getting help

Hewlett-Packard has application notes that describe how to configure Storage Mirroring with a variety of popular third-party applications. These application notes are available on the Storage Mirroring web site: http://h18006.www1.hp.com/products/storage/software/sm/index.html.).

Solution Overview

Pervasive SQL offers a high-performance, reliable, low-maintenance database for multiple operating systems. Storage Mirroring provides real-time enterprise data protection and replication. Storage Mirroring can be used to provide high availability for your Pervasive SQL server running on Windows 200x.

This document describes the steps necessary to configure Storage Mirroring to provide high availability for Windows 200x servers running Pervasive SQL 2000 or v8. These procedures allow a secondary server to assume the identity and role of a failed SQL server while maintaining the availability of SQL services with minimal disruption or data loss.

To complete these instructions, you will install Pervasive SQL and Storage Mirroring, and configure Storage Mirroring for replication and failover. Due to the complexities of these applications, this document is intended for network administrators with experience installing, configuring, and maintaining network applications including Storage Mirroring and Pervasive SQL.

Requirements

- Microsoft Windows 200x with the latest service pack
- Two licensed copies of Pervasive SQL 2000 or Pervasive SQL v8
- Two licensed copies of Storage Mirroring

NOTE: Storage Mirroring allows you to configure one target to monitor and failover for one or more source machines. In a one-to-one configuration, you will want to replicate your SQL data to the same location on the target so that failover is automatic. In a many-to-one configuration, each SQL data store will need to be replicated to a unique location and then renamed to the corresponding SQL directory on the source before failover occurs.

This application note focuses on a single SQL server being replicated to a single target.

Modifying the sample script files.

After you modify the sample scripts, save them with a new name to remove the `SAMPLE_` prefix. Copy the scripts to the directory where Storage Mirroring is installed.

The sample batch files provided are only examples. Because no two environments or configurations are exactly the same, you **MUST** modify the sample scripts in order to make the solution work in your environment.

Sample Implementation

This section describes an example of how to configure Storage Mirroring and Pervasive SQL. Use these procedures as a guideline for creating your own implementation.

Because no two environments or configurations are exactly the same, you will probably need to implement additional or different steps than what is documented here in order to make the solution work in your environment.

Install software on the source

1. Install Pervasive SQL and create a database on the source, if you have not already done so.

2. Record the drive and directory where SQL is installed. The default directory for Pervasive SQL is `<drive>:\PVSW`.

SQL Installation Drive and Directory: _____

NOTE: If you already have SQL installed and are not certain where the files are located, view the properties of a table to determine the location.

3. Install Storage Mirroring on the source machine using the installation defaults. See the *Storage Mirroring getting started* guide for further details.

Install and configure software on the target

1. Install Storage Mirroring on the target using the installation defaults. See the Storage Mirroring *Getting Started* guide for further details.
2. Install Pervasive SQL and create a database on the target using the same drive and directory specifications recorded in step 2 of "Install software on the source" on page 4.
3. Set the Pervasive SQL services to manual startup so that all SQL files are closed on the target and the Storage Mirroring source can replicate the changes.

- If you are using Pervasive SQL 2000, set the following services to manual startup:
  - Pervasive SQL 2000 Relational
  - Pervasive SQL 2000 Transactional
- If you are using Pervasive SQL v8, set the following services to manual startup:
  - Pervasive SQL Relational
  - Pervasive SQL Transactiona

NOTE: If a failure should occur, the failover and failback scripts that you will be creating will control the stopping and starting of the SQL services.

Configure and begin mirroring and replication

1. Select **Start, Programs, <GilsSansBody Bold>Double-Take, Management Console**.
2. Double-click your source machine to log on.
3. Right-click the source and select **Properties**.
4. On the Source tab, enable **Block Checksum All Files on a Difference Mirror** and click **OK**.
5. Right-click your source machine and select **New, Replication Set** and enter the desired name for the replication set.
6. Select the data you wish to protect. Most likely, this will only include the SQL data and log files. It is not necessary to replicate the application files (`.dll` and `.exe` files) since they already exist on the target machine. Select `<drive>:\PVSW` and any other directories (even if on different drives) that you may have created to store your SQL data files.
7. Right-click the replication set name and select **Save** to save the replication set.
8. Drag and drop the replication set onto the target. The Connection Manager dialog box opens.

9. The **Source Server**, **Target Server**, **Replication Set**, and **Route** fields will automatically be populated. If you have multiple IP addresses on your target, verify the **Route** field is set to the correct network path. (For detailed information on connecting a source and target, reference the Storage Mirroring *user's guide*.)
10. Select **One to One** to map the replication set data from the source to an identical volume/directory structure on the target.
11. Click **Connect** to start the mirror and replication processes

---

NOTE: If you start Pervasive SQL and mount the replicated databases on the target, or if the data on the target is otherwise modified, the data on the source and target will no longer match. If the updated data on the target is not needed, perform a full or difference with block checksum mirror from the source to the target. If the updated data on the target is needed, restore the data from the target to the source.

Configure failover and begin failure monitoring

1. If a failure occurs, you will want to have the SQL services start on the target machine automatically. To do this, create a batch file called `postover.bat` using the sample batch file below. Save the batch file to the same directory where your Storage Mirroring files are installed.

NOTE: After you modify the sample scripts, save them with a new name to remove the sample_ prefix. Copy the scripts to the directory where Storage Mirroring is installed.

The sample batch files provided are only examples. Because no two environments or configurations are exactly the same, you MUST modify the sample scripts in order to make the solution work in your environment.

---

sample_postover.bat

```
REM ***SAMPLE*** Pervasive SQL post-failover script

REM This sample batch file is provided as an example only. Because no two
REM environments or configurations are exactly the same, you MUST modify
REM this script in order to make the solution work in your environment.

REM By default, this script is configured for Pervasive SQL 2000.
REM If you are running Pervasive SQL v8, you will need to comment out the SQL 2000 commands
REM and use the v8 commands.

REM The following two commands are for Pervasive SQL 2000
net start "pervasive.sql 2000 (relational)"
net start "pervasive.sql 2000 (transactional)"

REM The following two commands are for Pervasive SQL v8
REM net start "pervasive.sql (relational)"
REM net start "pervasive.sql (transactional)"
```

2. After a failure is resolved, you will be ready to bring your source back online. At this time, you will want to stop the SQL services on the target automatically. To do this, create a batch file called `preback.bat` using the sample batch file below. Save the batch file to the same directory where your Storage Mirroring files are installed.

sample_preback.bat

```
REM ***SAMPLE*** Pervasive SQL pre-failback script

REM This sample batch file is provided as an example only. Because no two
REM environments or configurations are exactly the same, you MUST modify
REM this script in order to make the solution work in your environment.

REM By default, this script is configured for Pervasive SQL 2000.
REM If you are running Pervasive SQL v8, you will need to comment out the SQL 2000 commands
REM and use the v8 commands.

REM The following two commands are for Pervasive SQL 2000
net stop "pervasive.sql 2000 (relational)"
net stop "pervasive.sql 2000 (transactional)"

REM The following two commands are for Pervasive SQL v8
REM net stop "pervasive.sql (relational)"
REM net stop "pervasive.sql (transactional)"
```

3. Select **Start**, **Programs**, <GilsSansBody Bold>Double-Take, **Failover Control Center**.
4. Select the target machine from the list of available machines. If the target you need is not displayed, click **Add Target**, enter the machine name, and click **OK**.
5. To add a monitor for the selected target, click **Add Monitor**. Type the name of the source machine and click **OK**. The Monitor Settings window will open.
6. In the Monitor Settings window, mark the IP address that is going to failover.
7. Click **Scripts** and specify the scripts that were created earlier, using `postover.bat` for the target post-failover script and `preback.bat` for the target pre-failback script.

NOTE: The scripts are processed using the same account running the Double-Take service.

8. Click **OK** to go back to the Monitor Settings dialog box.
9. Click **OK** to begin monitoring the source machine.

In the event of a source machine failure, your target machine is now ready to stand in for your source. For detailed information on monitoring failover, see the *Storage Mirroring user's guide*.

Restoring your Pervasive SQL data

If your source experiences a failure, such as a power, network, or disk failure, your target machine will stand in for the source while you resolve the source machine issues. During the source machine downtime, data is updated on the target machine. When your source machine is ready to come back online, the data is no longer current and must be updated with the new data on the target machine.

1. Verify that your source machine is not connected to the network. If it is, disconnect it.
2. Resolve the source machine problem that caused the failure.

NOTE: If you must rebuild your hard drive, continue with step 3. If you do not need to rebuild your hard drive, continue with step 6 below.

3. Install Windows. Since your source machine is not connected to the network, go ahead and use the source's original name and IP address.
4. Install Storage Mirroring using the installation defaults.
5. Install Pervasive SQL using the same drive and directory settings recorded in step 2 of "Install software on the source" on page 4.

6. **Verify that SQL is not running on the source.** The SQL services must not be running at this time. Depending on the type of failure, your services may be set to manual startup but could still be running. **Stop your SQL services and set them to manual.**
7. Select **Start, Programs, <GilsSansBody Bold>Double-Take, Failover Control Center**.
8. Select the target machine that is currently standing in for the failed source.
9. Select the failed source and click Failback.

The pre-failback script entered during the failover configuration stops the SQL services on the target so that no additional changes can be made.

10. You will be prompted to determine if you want to continue monitoring the source server. Do not choose **Continue** or **Stop** at this time.
11. Connect the source machine to the network.
12. After the source is back online, select whether or not you want to continue monitoring this source machine (Continue or Stop).

NOTE: Verify that the Storage Mirroring connection on the source has been disconnected (right-click the connection in the Storage Mirroring Management Console and select **Disconnect**).

---

13. To begin the restoration process, open the Storage Mirroring Management Console and select **Tools, Restoration Manager**.

NOTE: You can also run the Storage Mirroring DTCL automated restoration script, which can be found in the Storage Mirroring *User's Guide*, to complete the remaining steps in this section.

---

14. Complete the appropriate fields as described below.

- **Original Source**—The name of the source machine where the data originally resided.
- **Restore From**—The name of the target machine that contains the replicated data.
- **Replication Set**—The name of the replication set to be restored.
- **Restore To**—The name of the machine where the data will be restored. This may or may not be the same as the original source machine.

15. Identify the correct drive mappings for the data and any other restoration options necessary. For detailed information on the restoration options, see the Storage Mirroring *User's Guide*.
16. Verify that the selections you have made are correct and click **Restore**. The restoration procedure time will vary depending on the amount of data that you have to restore.
17. After the restoration is complete, start the SQL services on the source machine and then reestablish the Storage Mirroring SQL replication set connection.

At this time, your data is restored back to your source machine, the source machine is again the primary SQL server, and, if you selected to continue failover monitoring, the target is available to stand in for the source in the event of a failure.